

**A METHOD AND APPARATUS FOR PERSISTENT ACCESS TO
WEB RESOURCES USING VARIABLE TIME-STAMPS**

5 This application is a continuation-in-part of U.S. Patent Application Serial No. 09/201,750, filed December 1, 1998, which is assigned to the assignee of the present invention and incorporated by reference herein.

Field of the Invention

10 The present invention relates to Internet resource access techniques, and more particularly, to a method and apparatus for ensuring persistent access to Internet resources.

Background of the Invention

15 The World Wide Web (the "Web") provides a dynamic way to present and distribute a vast amount of information. Anyone who is connected to the Internet and has a browser, such as Netscape Navigator Communicator™, commercially available from Netscape Communications Corporation of Mountain View, CA, can access information on the Web. The Web provides users with many media options and is becoming ubiquitously available in an expanding variety of personal electronic devices, far beyond its initial limited availability to 20 users via computer terminals. In addition, as display technologies continue to improve, the Web may ultimately replace traditional paper-based media altogether.

25 Paper-based media generally have an associated time stamp, and permit an easy determination of the information that was available at a given time. For example, a newspaper article can be cited as an authoritative reference, provided that the particular date of the newspaper publication is specified. Due to the dynamic nature of Web content, however, a Web document is generally not a reliable reference source. Currently, Web content cannot reliably be expected to be available in the same form and addressed by the same Uniform Resource Locator ("URL") at a future time. While some Web sites may provide access to some archived Web documents, the historical Web documents may not be accessed by users in a consistent and 30 predictable manner, if at all.

17 The Online Computer Library Center, Inc. ("OCLC"), a nonprofit computer library service and research organization, provides a software tool, referred to as OCLC PURL ("Persistent Uniform Resource Locator"), for managing Internet addresses and aliases for general Internet resources. A Persistent Uniform Resource Locator provides flexible naming and name resolution services for Internet resources to ensure reliable, long-term access to Internet resources with minimal maintenance. Generally, OCLC PURL assists Internet users in locating Web resources. As previously indicated, the Internet is constantly expanding and changing. Once a Uniform Resource Locator (URL) changes, all previous references to that URL become invalid, thereby preventing users from accessing the Internet resource. The management of these changes often becomes burdensome.

20 While a URL points directly to the location of an Internet resource, a PURL points to an intermediate resolution service, which translates the PURL into the actual URL. Once a Web resource has been registered with the OCLC and assigned a PURL, the Web resource may be accessed by means of the PURL. A PURL assigns a persistent name to a resource even if the location of the resource changes. In this manner, PURLs referenced in Web documents and other resources can remain viable over time without having to update the references each time the Web resource is moved. The PURL "forwarding" address maintained by OCLC, however, must be kept up-to-date. In other words, each time the document is moved, OCLC must be notified of the new address for the document.

20

Summary of the Invention

22 Generally, a method and apparatus are disclosed for providing persistent access to Web resources. According to an aspect of the invention, the Uniform Resource Locators ("URLs") that identify Web resources are optionally augmented to include a time stamp. The 25 time stamp can be specified in the Uniform Resource Locator ("URL") in any suitable format. In addition, variable time stamps can be utilized in accordance with the present invention to indicate a number of different dates, such as a date range or a recurring period of time.. According to one aspect of the invention, wildcard characters and date ranges can be used in the variable time-stamp to implement a variable time stamp when a user is not sure of the date for a specific web

resource or wishes to specify more than one precise date and time. The present invention allows the Web to be an organized and reliable reference source, much like paper-based media.

(u) A web browser and a web server are disclosed that accommodate a time stamp parameter and allow a user to refer to any Web address with a precise target time. The disclosed 5 Web browser can optionally include a mechanism to facilitate the specification of the desired date and time, or the user can manually append the time stamp to the URL indicated in the "Location" window of the browser. The persistent Web server (i) receives URLs containing a time stamp, a relative time-stamp or a variable time-stamp, (ii) retrieves the correct Web page(s) from an archive, and (iii) returns the desired page(s) or links to the client. The persistent Web 10 server interprets the extracted URL in accordance with the selected time stamp format. In addition, the persistent Web servers include a persistent archive for storing all of the versions of Web resources that will be persistently available to Web users. The present invention ensures that a time-stamped reference to any Web resource refers to the desired material. In this manner, anyone doing historical research on the Web can retrieve information that is valid in any period 15 of time.

Brief Description of the Drawings

FIG. 1 illustrates an Internet or World Wide Web ("Web") environment in accordance with the present invention;

20 FIG. 2 illustrates a directory structure that arranges the contents of a Web site chronologically;

FIG. 3 is a flow chart illustrating an exemplary archival process implemented by a Web site to reduce the redundancy of a persistent archive of FIG. 1;

25 FIGS. 4A and 4B illustrate the aliasing of the contents of a Web site that is identical to previously archived contents by the archival process of FIG. 3;

FIGS. 5A and 5B are a sample table of a DNS server database in accordance with the present invention; and

FIG. 6 is a flow chart describing an exemplary DNS server process performed by a DNS server to determine the IP address corresponding to a time-stamped URL.

Detailed Description

FIG. 1 illustrates a Web browser 100 in accordance with the present invention, that accesses information from one or more persistent Web servers 140, 150 over the Internet or World Wide Web ("Web") environment 130. The present invention provides persistent access to Web resources or electronic documents, including textual, audio, video or animation documents. According to a feature of the present invention, the Uniform Resource Locators ("URLs") that identify Web resources are augmented to include a time stamp. The Web browser 100 and persistent Web servers 140, 150 accommodate the additional time stamp parameter and allow a user to refer to any Web address with a precise target date. For example, the URL "http://cnn.com?time=2+2+1997" specifies the Web contents of the CNN home page on February 2, 1997. In this manner, the Web can be an organized and reliable reference source, much like paper-based media.

ILLUSTRATIVE TIME STAMP FORMATS

The time stamp can be included in the Uniform Resource Locator ("URL") in any suitable format, as would be apparent to a person of ordinary skill. For example, to refer to the web page, www.Lucent.com, as it existed on February 2, 1998, the URL can be represented as:

http://www.lucent.com?time=2+2+1998.

In a further variation, additional time granularity can be indicated by including the time-of-day in the URL. For example, the web page, www.Lucent.com, as it existed at 1:23 p.m. on February 2, 1998, the URL can be represented as:

http://www.lucent.com?time=13+23+2+2+1998.

Unless otherwise specified, the time zone is assumed to be the user's default time zone. The illustrative time stamp format described above is a Common Gateway Interface (CGI) search argument. Of course, the month, day and year (or other time units) can be expressed in any order. For a URL without a time stamp, the default value will be the most recent version.

In addition, relative time stamps can be utilized. For example, to refer to the web page, www.Lucent.com, as it existed yesterday, the URL can be represented as:

http://www.lucent.com?time=yesterday.

Furthermore, if an embedded hyperlink contains a relative time stamp, the relative time stamp is based on the current web page. Thus, if a current web page has a URL in the form:

5 http://www.lucent.com?time=1998_1_10

and this page contains an embedded hyperlink in the form:

5 <http://www.lucent.com?time=yesterday>

The browser and server will interpret and translate the URL as

10 http://www.lucent.com?rtime=1998_1_9.

Other relative time stamps can include time offsets from the time of the current web page, such as plus or minus a specified period of time. For example, "+10D" can indicate plus ten days to 10 the time of the currently viewed web page.

In order to refer to the previous or subsequent archived version of a document (relative to the time stamp of the current document), the URL can be represented using the labels "next_archive," or "previous_archive." In another variation, the first or most recent archived version of a document can be represented using the labels "first_archive," or "final_archive," 15 respectively. The server will search through the archive to find the required document. For example, if a currently viewed document has a URL in the form:

20 http://www.lucent.com/doc.html?rtime=1997_5_12

and there are different versions of the same document archived on 7/2/97, 6/1/97, 4/1/97 and 3/1/97. The following link can be used in the current document to refer to one of these archives:

20 [doc.html?rtime=next_archive](http://www.lucent.com/doc.html?rtime=next_archive)

These relative archive time stamps make moving between different version of the same documents more efficient.

A time base parameter can be used to specify the reference date for the relative timestamp. For example,

25 [timbase=1998_6_11&time=+100D](http://www.lucent.com?timbase=1998_6_11&time=+100D)

indicates the date that is 100 days after June 11, 1998. Similarly,

[timbase=1998_6_11&time=next_Monday](http://www.lucent.com?timbase=1998_6_11&time=next_Monday)

indicates the date that is the Monday after June 11, 1998.

According to another feature of the present invention, wildcard characters and date ranges in the time stamp can be used to implement a variable time stamp in a URL when a user is not sure of the date for a specific web resource. In this manner, the server can display a list of the specific web resources that match the time stamp pattern. In the illustrative 5 implementation, the following time stamp patterns are used:

TIME STAMP PATTERN	MEANING
*	wildcard character matching 0 or more digits
.	wildcard character matching one digit
m--n	from m to n
,	or
()	character grouping
---	date range specifier to specify a range between two dates (either absolute or relative dates)

The above time stamp patterns can be used in the time= or rtime= (relative time stamp) fields of the URL to specify, for example, an unknown year, month, day, hour, minute or second. For example,

EXAMPLE	MEANING
199(1,2)	1991 or 1992
199(0-9)	1990 through 1999
19(2-3)(4-5)	1924, 1925, 1934, or 1935
19(20-30, 88, 90)	1920 to 1930, 1988, or 1990
190.	1900 to 1909
19..	1900 to 1999
19*	19, 190 to 199, 1900 to 1999, ...

10

Thus, the time stamp patterns can be used to request a list of resources having a matching time stamp. For example

EXAMPLE	MEANING
http://www.a.com/res.html?rtime=1998_*	all res.html pages in 1998
http://www.a.com/res.html?rtime=1998_10_*	all res.html pages in October, 1998
http://www.a.com/res.html?rtime=*	all archived res.html pages
http://www.a.com/res.html?rtime=1998_10_(2,3)	all res.html pages on October 2 and 3, 1998

When a server receives a URL request containing a variable time stamp, the 15 server recognizes that the client is requesting a list of different versions of the same resource.

The server will search through all the archives to identify all matched resources and return an HTML page with hyperlinks pointing to all matched resources. Since some web resources might have many versions archived, the user can optionally specify how to present the links. Thus, according to a further feature of the invention, a "timeorder" parameter allows the user to specify 5 how to display the links corresponding to the matched resources. For example, timeorder=increase will present the links in increasing time order.

In addition, the links can be presented in a calendar-like format for easy navigation and selection. For example, if the links for the matching resources expand through several years, the links can be displayed in the following manner, with the month number 10 underlined to indicate existing versions of the matching resources:

1996	1	2	<u>3</u>	4	5	6	<u>7</u>	8	9	10	11	12
1997	1	2	3	4	<u>5</u>	6	7	<u>8</u>	9	10	11	12
1998	1	2	<u>3</u>	4	5	<u>6</u>	7	8	9	10	11	<u>12</u>

The URL corresponding to the link for March, 1998 would have the form http://www.a.com/res.html?rtime=1998_3_*&timeorder=calendar. The time order can also be specified in terms of units of time. For example, timeorder=+D means to list the links of the 15 matching resources in increasing day order, with the links corresponding to the first day of each month first, followed by links corresponding to the second day of each month and so on. Likewise, timeorder=+DY means to list the links of the matching resources in increasing day order, then in increasing year order.

Relative time stamps can be extended using the "*" wildcard at the end of the time 20 value. For example, if the current day is December 9, 1998, rtime=next_month* means rtime=1999_1_* (any day in January, 1999). Similarly, if the current day is December 9, 1998, rtime=+1y* means rtime=1999_12_9_* (any time on December 9, 1999).

In an alternate implementation, referred to herein as the "request-header scheme," the time stamp can be indicated as one of the HTTP request headers, such as:

25 Time-Stamp: June 9, 1998.

In another embodiment, referred to herein as the “special character scheme,” special characters can be utilized to indicate the inclusion of a time stamp in the URL, such as:

http://www.cnnfn.com;time=2+2+1998&timezone=server.

WEB BROWSER

5 The Web browser 100 may be embodied as a conventional browser, such as Microsoft Internet Explorer™ or Netscape Navigator™, as modified herein to incorporate the features and functions of the present invention. As discussed further below, the Web browser 100 only needs to incorporate a new options selection panel to permit the user to specify the desired date and time. In fact, a conventional Web browser 100 can be utilized, with the user 10 manually appending the time stamp to the URL indicated in the “Location” window of the browser 100.

In one implementation, the user has the option to turn the time stamp on or off. If the time stamp is activated, the browser 100 will change the URL accordingly before sending the URL out to the Web 130. Since there is no guarantee that the corresponding web server 140, 150 15 recognizes a time stamp, the document returned by the server 140, 150 might contain embedded hyperlinks that do not contain time stamps. Thus, in this situation, the web browser 100 can automatically convert the URL associated with an embedded hyperlink to add an appropriate time stamp when the user clicks on the hyperlink if the time stamp option is activated. The Web browser 100 should convert the URL in accordance with the selected time stamp format. In a 20 request-header-scheme implementation, the browser 100 should be modified to send the special request header (“Time-Stamp: June 9, 1998”). In addition, the HTML should be modified to include a new time stamp tag for any embedded hyperlink with a specific time stamp. For example, for a hyperlink such as:

Lucent Web Site

25 the HTML should be modified to indicate the time stamp of Feb. 2, 1998 as follows:

<TIMESTAMP TIME="2+2+1998" TIMEZONE=server></TIMESTAMP>Lucent Web Site.

PERSISTENT WEB SERVERS

The persistent Web servers 140, 150 may be embodied as conventional hardware and software, as modified herein to carry out the functions and operations described below. Specifically, the persistent Web servers 140, 150 need to know how to (i) receive URLs containing a time stamp or relative time-stamp, (ii) extract the time stamp, (iii) retrieve the Web page corresponding to the appropriate time-stamp, and (iv) return the requested page to the client. The persistent Web servers 140, 150 should interpret the extracted URL in accordance with the selected time stamp format.

If a version of the Web resource corresponding to the requested time does not exist, the present invention provides a version of the document stored time-wise in the vicinity of the requested target time. For example, the present invention may assume the Web resource has not changed from the previous archived version, and the version of the Web resource with the most recent time-stamp preceding the requested time is provided. Alternatively, the version of the Web resource with the next immediate time-stamp after the requested time is provided.

In addition, the persistent Web servers 140, 150 need to preserve all the information in their history of serving the Web. Thus, as shown in FIG. 1, each persistent Web server, such as the servers 140, 150, includes a persistent archive 145, 155, respectively, for storing all of the versions of Web resources that will be persistently available to Web users. The persistent archives 145, 155 may be embodied as any storage device, although a persistent (non-erasable) storage device such as CD-ROM, CR-R, WORM or DVD-ROM may be preferred.

For the persistent Web servers 140, 150 to support dated URLs, the persistent Web servers 140, 150 need to store all of their contents in a chronicle fashion to enable the retrieval of timely information. In one implementation, shown in FIG. 2, the persistent archives 145, 155 store the entire web site contents on permanent storage devices according to some sort of chronological directory structure. FIG. 2 shows a directory structure 200 that arranges the contents of the Web site chronologically. Thus, each leaf, such as the leaf 210, in the directory structure 200 corresponds to a dated URL. For example, a dated URL such as:

<http://www.nytimes.com?time=24+2+1998>

is conceptually equivalent to:

<http://www.nytimes.com/archive/1998/2/24/>.

Of course, storing the entire web site contents is inefficient in terms of storage usage. Many Web pages exhibit few, if any, changes from day to day. Thus, significant storage efficiencies can be achieved by simply removing redundancy in the archive. Once the 5 redundancy is removed, the storage requirement in addition to the regular web site storage is usually not very large.

FIG. 3 illustrates an archival process 300 for reducing the redundancy of the persistent archive 145, 155. All the files or subdirectories mentioned in the algorithm are under the archive subdirectory 220 of the illustrative directory structure 200 of FIG. 2. As shown in 10 FIG. 3, the archival process 300 initially performs a test during step 310 for each subdirectory, such as subdirectory A, to determine whether there exists a subdirectory B that is created earlier and has identical contents as subdirectory A. If it is determined during step 310 that there is no subdirectory B created earlier and having identical contents as subdirectory A, then it is not possible to reduce the redundancy on the subdirectory level of the persistent archive 145, 155 and 15 program control proceeds to step 330.

If, however, it is determined during step 310 that there exists a subdirectory B that is created earlier and has identical contents as subdirectory A, then subdirectory A becomes an alias during step 320 pointing to subdirectory B. For example, as shown in FIG. 4A, if the current contents of a Web site is identical to the contents of the previous day, an alias is created 20 for today pointing to yesterday's subdirectory. Likewise, as shown in FIG. 4B, if the current month's content is the same as the contents of the previous month, an alias is created for this month pointing to last month's subdirectory.

Thereafter, a test is performed during step 330 for each file, such as file A, to determine whether there exists a file B that is created earlier and has identical contents as file A. 25 If it is determined during step 330 that there is no file B created earlier and having identical contents as file A, then it is not possible to reduce the redundancy of the persistent archive 145, 155 on the file level. Thus, program control terminates during step 350.

If, however, it is determined during step 330 that there exists a file B that is created earlier and has identical contents as file A, then file A becomes an alias during step 340 pointing to file B. Thereafter, program control terminates during step 350.

The archival process 300 may be impractical, since it needs to search for match files or directories. The run time increases exponentially with the number of entities in the archive. Many sub-optimal solutions are possible, as would be apparent to a person of ordinary skill in the art. A very simple solution is just checking what you want to archive today against the most recently added archive (like yesterday's contents). Since most of the web sites only differ from their previous archived ones slightly, this approach is quite reasonable. This approach is similar to the well-known incremental backup of a file system.

If a Web server is not persistent, it should only have minimal impact. In one embodiment, if a request includes a time stamp that is not recognized by a Web server, the server should deliver the most recent version of the requested Web resource.

Another way to reduce storage requirements of the persistent archive is to make the Web server smarter in terms of searching the correct archived data. For example, persistent storage of a web resource can be limited to versions that have some difference relative to previously saved versions of the web resource. For example, if an illustrative archive contains the following five different versions of a web resource: 6/4/1996, 6/12/1996, 3/23/1997, 2/1/1998 and 2/3/1998, the web server assumes that if the requested date does not equal any of the archived versions, then the requested date is identical to the version with the closest earlier date. In addition, a special symbolic link (or alias on MacOS, short cut on MS Windows) can be used in a directory to represent where to look for files or directories that are not found under the current directory. In this manner, only the changed parts are stored under appropriate directories. All the unchanged data can be referred through a chain of such special links.

25 DNS SERVER

The domain name server (DNS) may be embodied as conventional hardware and software, as modified herein to carry out the functions and operations described below. Conventional DNS servers will reject any domain name reference which is not in the DNS database. One benefit of dated URL in accordance with the present invention is that it can be

used to refer to historical Web resources. For example, if company A is merged into company B, all the web pages referred through “www.A.com” may no longer be valid. For users who want to access some documents from company A, they need to change all the reference to some place in company B’s web site.

5 The historical information of company A can still be accessed if the DNS server does not reject the name reference, but instead consults an archive service company that knows where the historical information of company A is located. The DNS server itself can also store some historical data to resolve the name to IP address process faster. FIGS. 5A and 5B provide examples of data stored in a DNS server database before and after the merger of companies A
10 and B, respectively. As shown in FIG. 5B, if a user wants to find www.A.com after the merger, the DNS server has enough information to redirect the user’s request to a new IP address associated with company B. The dates listed in the database are the valid periods for the corresponding domain name. Thus, a dated domain name reference like “www.A.com 2/2/1999” is invalid, while “www.A.com 2/2/1992” is valid.

15 FIG. 6 illustrates a DNS server process 600 in accordance with the present invention. As shown in FIG. 6, the DNS server process 600 initially receives a domain name request during step 610. A test is performed during step 620 to determine if the domain name request is dated. If it is determined during step 620 that the domain name request is not dated, the regular name searching result is returned during step 630.

20 If, however, it is determined during step 620 that the domain name request is dated, the DNS server process 600 searches the DNS database for the domain name with the date constraint during step 640. A further test is performed during step 650 to determine if the dated domain name is found. If it is determined during step 650 that the dated domain name is not found, then the DNS server consults with an archive service company during step 660 for further
25 searching before program control proceeds to step 670.

If, however, it is determined during step 650 that the dated domain name is not found, then the searching result and indication, if redirect, are returned during step 670, before program control terminates.

After the domain name is resolved by the DNS server, the Web browser 100 needs to send the request to the web server 140, 150 according to what is returned from the DNS server. For example, a request from the user for the following URL, "http://www.A.com?time=2+2+1992," will cause the browser 100 to send a domain name 5 resolving request to the DNS server in a format such as "www.A.com 2/2/1992." Since company A is now part of company B, the results will look like "123.2.3.222 redirect." The Web browser 100 now has the IP address of the server and also knows it is a redirect one. Thus, the Web browser 100 will effectively send a request to the Web server 140, 150 of Company B in a form such as "http://123.2.3.222?http://www.A.com&time=2+2+1992." The Web server 140, 10 150 of Company B will know how to map this old address of company A's to the appropriate place and get the correct information.

One side benefit of this new DNS server is that some names can be reused once they are history. For example, another company named Company A can utilize the www.A.com domain name after a predefined period, by updating the DNS database with the following entries:

URL	STARTING DATE	ENDING DATE	IP ADDRESS	REDIRECTED URL
WWW.A.COM	13/2/1980	1/5/1998	123.2.3.222	WWW.B.COM
WWW.A.COM	1/1/2000	PRESENT	234.2.2.12	--
WWW.B.COM	23/1/1985	PRESENT	123.2.3.222	--

15

In this manner, domain names can be reused without wasting them forever.

PERSISTENT ACCESS TO DYNAMIC WEB CONTENT

The Web is now full of dynamic content, including real time video, for example, from a WebCam, and audio streams, for example, from a WebCast event, as well as Java, 20 Javascript or Active-X enabled web pages. Depending on the application, it may not be necessary or desirable to archive the continuous dynamic content. In addition, it may not be necessary or desirable to archive all the advertising portions of a web pages.

For chronological data, such as bank or stock broker transactions, it is easy to extract part of the record for a given time stamp restriction. For example, to check the account 25 balance of a give date, the browser 100 can send a request in the form:

<http://bank.com/?id=12345&time=2+3+1998>.

The server 140, 150 only needs to retrieve or recalculate the data up to March 2, 1998 and return the results. Since all the transactions in such application environments have time stamps anyway, it is straightforward to add this function to the service.

5 For real time contents, the only restriction in appending a time stamp is the storage requirement. If a lot of storage space is available compared to the amount of information to be archived, the Web site administrator can choose to archive the real time contents or to archive some of them such as one day, one week or one year's worth of data.

10 For dynamically created advertisements, the Web site administrator must decide whether it is reasonable to 'reshow' the old advertisement (for some special reason) or whether the old advertisement can be replaced with a new, up-to-date commercial which is not relevant to the 'real' archived web contents.

15 For a dynamically executed script on a Web page, there are two possible situations. If the script is not related to any time function, then the script can be archived safely. If, however, the script will change its behavior according to some time function, such as a clock or calendar display, then the script needs to be modified to deal with dated URL, requiring modifications to Java, Javascript or Active-X standards, as appropriate. Web pages generated dynamically by CGI applications are treated in a similar manner.

20 It is to be understood that the embodiments and variations shown and described herein are merely illustrative of the principles of this invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention.